



UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
10/777,508	02/12/2004	Alfredo Aldereguia	RPS9 2003 0207 US1	5244
56102	7590	11/26/2008		
IBM (RPS-BLF)				
c/o BIGGERS & OHANIAN, LLP				
P.O. BOX 1469				
AUSTIN, TX 78767-1469				
EXAMINER				
AHMED, SALMAN				
ART UNIT		PAPER NUMBER		
2419				
MAIL DATE		DELIVERY MODE		
11/26/2008		PAPER		

Please find below and/or attached an Office communication concerning this application or proceeding.

The time period for reply, if any, is set in the attached communication.



UNITED STATES PATENT AND TRADEMARK OFFICE

Commissioner for Patents
United States Patent and Trademark Office
P.O. Box 1450
Alexandria, VA 22313-1450
www.uspto.gov

**BEFORE THE BOARD OF PATENT APPEALS
AND INTERFERENCES**

Application Number: 10/777,508
Filing Date: February 12, 2004
Appellant(s): ALDEREGUIA ET AL.

H. Artoush Ohanian
For Appellant

EXAMINER'S ANSWER

This is in response to the appeal brief filed 8/26/2008 appealing from the Office action mailed 4/16/2008.

(1) Real Party in Interest

A statement identifying by name the real party in interest is contained in the brief.

(2) Related Appeals and Interferences

The examiner is not aware of any related appeals, interferences, or judicial proceedings which will directly affect or be directly affected by or have a bearing on the Board's decision in the pending appeal.

(3) Status of Claims

The statement of the status of claims contained in the brief is correct.

(4) Status of Amendments After Final

No amendment after final has been filed.

(5) Summary of Claimed Subject Matter

The summary of claimed subject matter contained in the brief is correct.

(6) Grounds of Rejection to be Reviewed on Appeal

The appellant's statement of the grounds of rejection to be reviewed on appeal is correct.

(7) Claims Appendix

The copy of the appealed claims contained in the Appendix to the brief is correct.

(8) Evidence Relied Upon

2003/0158940	Leigh et al.	02-2002
2004/0057448	Nakamura	09-2003
6286073	Vegter	01-1998

(9) Grounds of Rejection

The following ground(s) of rejection are applicable to the appealed claims:

Claim Rejections - 35 USC § 102

1. Claims 1-2, 8-10, and 15-17 are rejected under 35 U.S.C. 102(b) as being anticipated by Leigh (US2003/0158940).

Regarding claim 1, Leigh discloses method for integrated load balancing among peer servers (**see paragraph 7**) comprising:

a first set of central processing units (see paragraph 29 CPUs);

a first system memory accessible to the first set of processors (**see paragraph 29 processors**);

scalability logic (**see paragraph 29 scalability logic is not further specified therefore it is broadly interpreted as load balancing algorithms**) to connect the data processing system to a second data processing system, having a second set of processors and a second system memory, to form a scaled system (**see paragraph 29 and figure 1c**);

a set of scalability ports (**see paragraph 40 the configuration shown is only one possible arrangement of connections among the ports; many other permutations are possible. Not all of the possible connection configurations are optimal for a chosen load balancing method**) connected to the scalability logic to receive scalability cables (see paragraph 40 interconnecting cables among the ILBs)

connecting the first system to the second system (see paragraph 29 and figure 1a-d);
and

system management to cause each of the system's scalability ports (see paragraph 40 the configuration shown is only one possible arrangement of connections among the ports; many other permutations are possible. Not all of the possible connection configurations are optimal for a chosen load balancing method) to issue an identifiable signal (see paragraph 30 request) and further configured to detect the reception of an identifiable signal (see paragraph 30 ILB 50 acknowledges the load shedding request from ILB 10), sent by another system, by any of the scalability ports (see paragraph 40 the configuration shown is only one possible arrangement of connections among the ports; many other permutations are possible. Not all of the possible connection configurations are optimal for a chosen load balancing method) and to report the reception of the signal to a system management of the second system to determine which ports of the two systems are connected by the cable (see paragraph 47 and figure 1 a-d each port knows the ID of the port of ports to which it is directly connected and see paragraph 40 interconnecting cables among ILBs).

Regarding claim 2, Leigh teaches the system management includes a service processor connected to the system via an adapter card and wherein the service processor is connected to a service processor of other systems via a network medium (see paragraph 49 figure 1a-d and figure 3 zone A-D ref311-4 and ref306 external network).

Regarding claim 8, Leigh teaches a method of determining scalability cabling between at least two scalable data processing systems, comprising:

driving an identifiable signal on a first scalability port (**see paragraph 40 the configuration shown is only one possible arrangement of connections among the ports; many other permutations are possible. Not all of the possible connection configurations are optimal for a chosen load balancing method**) of a first system (**see paragraph 30 request**);

responsive to receiving the identifiable signal by a second system, determining which scalability port (**see paragraph 40 the configuration shown is only one possible arrangement of connections among the ports; many other permutations are possible. Not all of the possible connection configurations are optimal for a chosen load balancing method**) of the second system received the distinctive signal (**see paragraph 30 acknowledge**);

informing the first system of the reception of the distinctive signal (**see paragraph 30 acknowledge**) by the determined scalability port (**see paragraph 40 the configuration shown is only one possible arrangement of connections among the ports; many other permutations are possible. Not all of the possible connection configurations are optimal for a chosen load balancing method**) of the second system and recording the first scalability port (**see paragraph 40 the configuration shown is only one possible arrangement of connections among the ports; many other permutations are possible. Not all of the possible connection configurations are optimal for a chosen load balancing method**) of the first system

and the scalability port of the second system as being connected by a scalability cable (see paragraph 30 ILB 0 provides ILB 10 with a list of ILBs that can accept the load and with the load-shedding conditions).

Regarding claim 9, Leigh teaches further comprising detecting a timeout by the first system and, responsive thereto, identifying the first scalability port (see paragraph 40 the configuration shown is only one possible arrangement of connections among the ports; many other permutations are possible. Not all of the possible connection configurations are optimal for a chosen load balancing method) as being unconnected (see paragraph 47 port 24 know it is not connected and figure 1b ref 24).

Regarding claim 10, Leigh teaches further comprising, iterating the sequence of claim 8, until all scalability ports (see paragraph 40 the configuration shown is only one possible arrangement of connections among the ports; many other permutations are possible. Not all of the possible connection configurations are optimal for a chosen load balancing method) have been accounted for (see paragraph 47 and figure 1a-d and paragraph 30 ILB 0 provides ILB 10 with a list of ILBs that can accept the load and with the load-shedding conditions).

Regarding claims 15-17, Leigh teaches programming (see paragraph 42) and disclose all the limitations as discussed in the rejection of claims 8-10 and are therefore claims 15-17 are rejected using the same rationales.

Claim Rejections - 35 USC § 103

2. Claims 3, 6-7, 11-12, 14, 18-19, and 21 are rejected under 35 U.S.C. 103(a) as being unpatentable over Leigh in view of Nakamura (US2004/0057448).

Regarding claims 3 and 14, Leigh teaches

the system management causes a scalability port **(see paragraph 40 the configuration shown is only one possible arrangement of connections among the ports; many other permutations are possible. Not all of the possible connection configurations are optimal for a chosen load balancing method)** to issue an identifiable signal

disclose all the subject matter of the claimed invention with the exception of

by causing the assertion of a bit in a register corresponding to the set of scalability ports;

determining the scalability port that received the signal includes reading the bits in a register associated with scalability port.

Nakamura from the same or similar fields of endeavor teaches the use of the communication ports of those nodes are assigned port numbers to identify these ports, the 1394 interface of each node automatically recognizes a new connection configuration, and assignment of node IDs **(see Nakamura paragraph 122 and paragraphs 179 and 180 where node ID is being assigned and paragraph 152 communication ports of those nodes are assigned port number to identify these ports and paragraph 155 nodes declare parent-child relationships among their communication port, therefore when assigning node ID, ports are also assigned**

with their relationship to other nodes and figure 6 ref602 node ID# and figure 14, fig. 15 fig. 16), and upon receiving the self ID packet, each node can recognize the node numbers assigned to the respective nodes and can detect a node number assigned to itself (see Nakamura paragraph 187).

Thus, it would have been obvious to one of ordinary skill in the art at the time of the invention to use the assigning node ID and identify communication ports as taught by Nakamura in the integrated load balancing among peer server of Leigh in order to enhance system efficiency.

Regarding claim 6, Leigh teaches identifying the corresponding scalability port (see paragraph 40 the configuration shown is only one possible arrangement of connections among the ports; many other permutations are possible. Not all of the possible connection configurations are optimal for a chosen load balancing method) open (see Leigh paragraph 47 and figure 1b port 24) and disclose all the subject matter of the claimed invention with the exception of the system management includes means for determining a timeout condition following assertion of a bit.

Nakamura from the same or similar fields of endeavor teaches the use of the communication ports of those nodes are assigned port numbers to identify these ports, the 1394 interface of each node automatically recognizes a new connection configuration, and assignment of node IDs (see Nakamura paragraph 122 and paragraph 152-162 and figure 6 ref 602 node ID# and figure 14, fig. 15 fig. 16).

Thus, it would have been obvious to one of ordinary skill in the art at the time of the invention to use the assigning node ID and identify communication ports as taught

by Nakamura in the integrated load balancing among peer server of Leigh in order to enhance system efficiency.

Regarding claim 7, Leigh disclose all the subject matter of the claimed invention with the exception of the system management further includes code means for using the scalability information to generate a graphical image of scalability interconnections.

Nakamura from the same or similar fields of endeavor teaches the use of display function (**see Nakamura paragraph 236 and 237**).

Thus, it would have been obvious to one of ordinary skill in the art at the time of the invention to use the display function as taught by Nakamura in the integrated load balancing among peer server of Leigh in order to provide enhance efficiency in system monitoring.

Regarding claims 11-12, 18-19, and 21, Leigh and Nakamura disclose all the limitations as discussed in the rejection of claims 3, 7, and 14 and are therefore claims 11-12, 18-19, and 21 are rejected using the same rationales.

3. Claims 4 is rejected under 35 U.S.C. 103(a) as being unpatentable over Leigh and Nakamura further in view of Vegter (US.6286073).

Regarding claim 4 and 5, Leigh and Nakamura disclose all the subject matter of the claimed invention with the exception of

the scalability port register is implemented in a programmable logic device

the system management further includes controller logic connected to the service processor via a dedicated serial connection and connected to the programmable logic device via an 12C bus.

Vegter from the same or similar fields of endeavor teaches the use of 12C control block may be a programmable logic array (see Vegter col. 3 lines 21-42).

Thus, it would have been obvious to one of ordinary skill in the art at the time of the invention to use the programmable logic array as taught by Vegter in the modified integrated load balancing among peer server of Leigh and Nakamura in order to provide enhance efficiency in system monitoring solution for interfacing integrated circuits to one another, when it is necessary to interface an external 12C device with a personal computer, additional hardware is typically needed which is expensive, cumbersome to use and not versatile. Thus, the integrated circuit art has not heretofore created a solution for simple and effective 12C interfacing between a personal computer and any 12C device (see Vegter col. 1 lines 46-54).

4. Claims 5, 13, and 20 are rejected under 35 U.S.C. 103(a) as being unpatentable over Leigh in view of Vegter (US.6286073).

Regarding claim 5, Leigh disclose all the subject matter of the claimed invention with the exception of

the system management further includes controller logic connected to the service processor via a dedicated serial connection and connected to the programmable logic device via an 12C bus.

Vegter from the same or similar fields of endeavor teaches the use of 12C control block may be a programmable logic array (see Vegter col. 3 lines 21-42), and it is well known in the art to use the service processor via a dedicated serial connection and programmable logic device via an 12C bus.

Thus, it would have been obvious to one of ordinary skill in the art at the time of the invention to use the 12C control block may be a programmable logic array as taught by Vegter in the modified integrated load balancing among peer server of Leigh and Nakamura in order to provide enhance efficiency in system monitoring solution for interfacing integrated circuits to one another, when it is necessary to interface an external 12C device with a personal computer, additional hardware is typically needed which is expensive, cumbersome to use and not versatile. Thus, the integrated circuit art has not heretofore created a solution for simple and effective 12C interfacing between a personal computer and any 12C device (see Vegter col. 1 lines 46-54).

Regarding claims 13 and 20, Leigh, Nakamura, and Vegter disclose all the limitations as discussed in the rejection of claims 5 and are therefore claims 13 and 20 are rejected using the same rationales.

(10) Response to Argument

Appellant has itemized the arguments traversing the rejections of the appealed claims, each of which will be treated in turn.

35 USC § 102(b) - Claim 1

Appellant argues (page 6, paragraph 3) that *Leigh Does Not Disclose Scalability Logic To Connect The Data Processing System To A Second Data Processing System*,

Having A Second Set Of Processors And A Second System Memory, To Form A Scaled System.

However, Examiner respectfully disagrees with the Appellant's assertion. Leigh does indeed teach the cited limitations. Specifically, Leigh teaches scalability logic (**see paragraph 29 scalability logic is not further specified therefore it is broadly interpreted as load balancing algorithms**) to connect the data processing system to a second data processing system, having a second set of processors and a second system memory, to form a scaled system (**see paragraph 29 and figure 1c**, the four-ILB (i.e. data processing systems) embodiment shown in FIG. 1(c) can be used as an example to describe the zone-based load balancing method. Assume that ILBs 10 and 20 belong to zone-1 and ILBs 30 and 50 belong to zone-2 and that ILBs 10 and 50 are the primary master ILBs for their respective zones. Simpler load-balancing algorithms, such as a basic round-robin method, may be used when the slave ILBs do not overlap across the zones. In the example of FIG. 1(c), slave ILB 20 is in zone-1 only and slave ILB 30 is in zone-2 only, although it is possible to assign ILB 20 and ILB 30 to be in both zone-1 and zone-2. A primary master ILB, e.g., ILB 10, can collaborate with another primary master ILB, e.g., ILB 50, for load sharing. Each primary master ILB has its own zone as its primary zone and the collaborating zone as its secondary zone. A primary master ILB may decide to cross the zone boundaries and send some of its load to the servers in another zone based on based on factors such as its host server's workload index, network data traffic level, server health, and ILB health (i.e. scalability logic to connect the data processing system to a second data processing system). "Server

health" constitutes functional statuses on the server host's critical subsystem components, such as CPU(s), cache memory, system memory, and disks (i.e. data processing system, having a second set of processors and a second system memory) "ILB health" constitutes functional statuses on the ILB's components, such as processors, memory buffers, ASICs, and FPGAs).

Appellant argues (page 7, paragraph 2) that *Leigh's load balancing with ILBs does not disclose scalability logic connecting one data processing system to a second data processing system as claimed here. It is clear to a person of skill in the art that load balancing as disclosed in Leigh simply has nothing to do with scalability logic for connections among systems. In Leigh, the servers among whom load balancing is implemented are expressly already connected in a network, therefore requiring no scalability logic for connections: as claimed in the present application. On the face of Leigh, therefore, there would be no reason for a person of skill in the art to expect to derive, anything regarding scalability for connections among systems from Leigh, which is clearly about load balancing and not about scalability.*

However, Examiner respectfully disagrees with the Appellant's assertion and interpretation. A quick search of the term "Scalable" in World Wide Web (internet) provides some of the following results (NOTE: Such results are not used here to provide a new ground of rejections, but, rather to show the contrast between "Scalability" and "Load-Balancing" and to show the understanding of one of ordinary skilled in the art):

<http://www.businessdictionary.com/definition/scalability.html>

The screenshot shows the BusinessDictionary website. The header includes the site logo, navigation links (Home, Browse by Subject, Most Popular, Term of the Day), and social media links (Add to Google, Add to Facebook). A search bar is present with the text "Search for a term: Enter a word or phrase...". Below the search bar, there are several advertisements. The main content area displays the definition of "scalable".

scalable

Definition 1

Facility, plant, or unit whose size, performance, or number of users can be increased on demand without a penalty in cost or functionality.

Definition 2

System designed to handle proportionally very small to very large usage and service levels almost instantly, and with no significant drop in cost effectiveness, functionality, performance, or reliability. Scalable systems employ technologies such as automatic load balancing, clustering, and parallel processing.

Experience More in HD
Comcast Has Over 1,000 HD Choices See the Difference For Yourself
www.comcast.com/HD

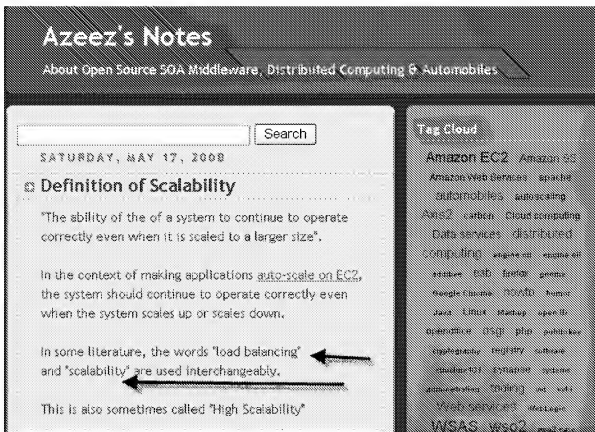
Software Scalability
Need help finding and fixing scalability problems?
www.cymatic.com

Maximize Usable Storage
FREE white paper "Understanding the Usable versus Raw Storage Capacity"
www.terracala.com

Top Grid Software
Get Grid MP 5.6 and save millions Download Free ROI white paper
www.ithread.com

Ads by Google

<http://afkham.org/2008/05/definition-of-scalability.html>



As such Examiner respectfully disagrees with Appellant's assertion that *a person of skill in the art that load balancing as disclosed in Leigh simply has nothing to do with scalability logic for connections among systems.*

In response to Appellant's argument that *In Leigh, the servers among whom load balancing is implemented are expressly already connected in a network, therefore requiring no scalability logic for connections*, Examiner respectfully points out that Leigh further teaches in paragraph 0057, when the workload index is consistently low for all the servers, some servers may be selected to be in "sleep" mode (i.e. not connected state) to reduce facility power consumption. If the primary master ILB detects a low workload index for all the servers, it sends a roll-call-sleep command packet to the ILBs

attached to the servers selected to be in "sleep" mode (i.e. not connected state). After the primary master ILB sends the roll-call-sleep packet to a server, it will not distribute any loads to that server. The primary master ILB adjusts the workload index ratio among the remaining servers accordingly, as described in the adaptive load balancing method. Upon receiving a roll-call-sleep command packet, a server will respond to the primary master ILB with a roll-call-sleep-wait packet. After the server finishes processing its outstanding assigned loads, it will send a roll-call-asleep packet, via its ILB, to the primary master ILB. When the primary master ILB receives the roll-call-asleep packet, it will send a roll-call-asleep-acknowledge packet to the server and the server will go into a sleep mode (i.e. not connected state). However, the ILB of a sleeping server stays awake to route loads passing through the internal network, when necessary, and to communicate with the primary master ILB to later wake up the sleeping server. When the workload index average of the remaining servers exceeds a certain workload threshold, e.g., 75% of their peak workload index, the primary master ILB can wake up (i.e. connected state) one or more servers that are in sleep mode. To wake up (i.e. connected state) a server that was put to sleep, the primary master ILB sends a roll-call-wakeup packet to the ILB that is attached to the sleeping server (i.e. not connected state). The ILB then wakes up the server (i.e. connected state) and responds to the primary master ILB with a roll-call-awaken packet when the server is ready to participate again, i.e., ready to receive loads to be processed (i.e. connected state). The primary master ILB then readjusts the workload index ratio among all the participating servers and continues with the load distribution (i.e. Scalability). The ability

to place servers in sleep mode (i.e. not connected state) and the ability to wake them up (i.e. connected state) in a short time to participate in load balancing enables scaling the computer resources while efficiently using the facility power.

As such Examiner respectfully disagrees with Appellant's assertion that *on the face of Leigh, therefore, there would be no reason for a person of skill in the art to expect to derive, anything regarding scalability for connections among systems from Leigh, which is clearly about load balancing and not about scalability.*

Appellant argues (page 7, last paragraph) that *Leigh does not disclose a "scaled system" as claimed in the present application. In the original specification at page 4, lines 4-7. Appellants define a "scaled system": as a system in which two or more symmetric multiprocessor ("SMP") systems are interconnected to form a larger multiprocessor system capable of functioning under a single operating system image. Neither at the references point cited by the examiner nor anywhere else in Leigh does Leigh disclose a scaled system capable of functioning under a single operating system image. In fact, the terms "image," "operating system," and "operating system image" do not appear at any point in Leigh - not even once. For these reasons, a person of skill in the art would understand a "scaled system." as the term in used in the present claims read in view of the specification, to be interconnected SMPs capable of running under a single operating system image. Moreover, it is clear to a person of skill in the art that all the servers described in Leigh are assumed to be operating their own separate operating systems.*

However, Examiner respectfully disagrees with the Appellant's assertion. In response to applicant's argument that the references fail to show certain features of applicant's invention, it is noted that the features upon which applicant relies (i.e., *"a system in which two or more symmetric multiprocessor ("SMP") systems are interconnected to form a larger multiprocessor system capable of functioning under a single operating system image"*, *"a scaled system capable of functioning under a single operating system image"*, *"image,"* *"operating system,"* and *"operating system image"*) are not recited in the rejected claim(s). Although the claims are interpreted in light of the specification, limitations from the specification are not read into the claims. See *In re Van Geuns*, 988 F.2d 1181, 26 USPQ2d 1057 (Fed. Cir. 1993).

As such Examiner respectfully disagrees with Appellant's assertion that *Leigh* cannot possibly disclose scalability logic to connect the data processing system to a second data processing system, having a second set of processors and a second system memory, to form a 'scaled system' as claimed here.

Appellant argues (page 8, last paragraph) that *Leigh* does not disclose the "data processing system" claimed in the present application, in the original specification at page 4, lines 11 - 12, Appellants define a "data processing system" as including a Processor Scalability and Cache control unit which controls access to tile scalability ports in the data processing system. A person of skill in the art therefore would understand a "data processing system," as the term is used in the present claims read in view of the specification, to include a Processor Scalability and Cache control unit which controls access to scalability ports. *Leigh* does not disclose a data processing system

containing scalability ports or a data processing system containing a Processor Scalability and Cache control unit which controls access to the scalability ports.

However, Examiner respectfully disagrees with the Appellant's assertion. In response to applicant's argument that the references fail to show certain features of applicant's invention, it is noted that the features upon which applicant relies (i.e., "*Processor Scalability and Cache control unit which controls access to tile scalability ports in the data processing system*") are not recited in the rejected claim(s). Although the claims are interpreted in light of the specification, limitations from the specification are not read into the claims. See *In re Van Geuns*, 988 F.2d 1181, 26 USPQ2d 1057 (Fed. Cir. 1993).

As such Examiner respectfully disagrees with Appellant's assertion that *because Leigh does not disclose each and every element and limitation of Appellants' claims, Leigh does not anticipate Appellants' claims and the rejections under 35 U.S.C. § 102 should be withdrawn.*

Appellant argues (page 9, first paragraph) that *Leigh Does Not Disclose A Set Of Scalability Ports Connected To The Scalability Logic To Receive Scalability Cables Connecting The First System To The Second System.*

However, Examiner respectfully disagrees with the Appellant's assertion. Leigh does indeed teach the cited limitations. Specifically, Leigh teaches a set of scalability ports **(see paragraph 40, the configuration shown is only one possible arrangement of connections among the ports; many other permutations are possible. Not all of the possible connection configurations are optimal for a**

chosen load balancing method. Leigh in paragraph 40 states, the use of 3-port ILB modules (i.e. A Set Of Scalability Ports) allows the creation (i.e. connoting) of load-balancing clusters (i.e. Scalability Logic) at low cost, since no additional hardware, such as a network switch, is needed. This is important especially when there are only a few servers to be load balanced and an ILB module in each host server can provide the load-balanced cluster solution. When there are only a few servers, the installation of the interconnecting cables (i.e. Scalability Cables) among the ILBs is not a significant issue. FIG. 1(d) shows an example of how five ILBs 10, 20, 30, 40, and 50 can be interconnected among themselves and to two external networks via network segments 80 and 90 (i.e. A Set Of Scalability Ports Connected To The Scalability Logic To Receive Scalability Cables Connecting The First System To The Second System). Six inter-ILB connections 83, 84, 85, 86, 87, and 88 are used to interconnect the five 3-port ILBs (i.e. A Set Of Scalability Ports Connected To The Scalability Logic To Receive Scalability Cables Connecting The First System To The Second System). The configuration shown is only one possible arrangement of connections among the ports; many other permutations are possible. Not all of the possible connection configurations are optimal for a chosen load balancing method. Similarly, in the embodiment shown in FIG. 2, twelve 3-port ILBs 202 through 224 are interconnected by sixteen connections and four primary master ILBs 202, 208, 214, and 220 are connected to four external networks via connections 230, 240, 250, and 260 (i.e. A Set Of Scalability Ports Connected To The Scalability Logic To Receive Scalability Cables Connecting The First System To The Second System)) connected to the scalability logic to receive scalability

cables (see paragraph 40 interconnecting cables among the ILBs) connecting the first system to the second system (see paragraph 29 and figure 1a-d).

Appellant argues (page 10, first paragraph) that *Leigh does not disclose scalability cables, scalability ports, or scalability logic. Scalability cables are distinct from network cables, such as Ethernet cables, that are used for connecting machines within a local area network. See, Appellants' original specification at page 1. lines 16-18 (It is worth noting here that the scalability cables referred to in this disclosure are distinct from the network cables (such as Ethernet cables) that connect systems in a local area network). Neither at the reference point cited by the Examiner nor anywhere else in Leigh does Leigh disclose scalability cables. Leigh merely discloses that servers are connected through the network ports of Leigh's internal load balancers, without making any disclosure regarding the cables that are used to connect Leigh's servers.*

However, Examiner respectfully disagrees with the Appellant's assertion. In response to applicant's argument that the references fail to show certain features of applicant's invention, it is noted that the features upon which applicant relies (i.e., *"the scalability cables referred to in this disclosure are distinct from the network cables (such as Ethernet cables) that connect systems in a local area network"*) are not recited in the rejected claim(s). Although the claims are interpreted in light of the specification, limitations from the specification are not read into the claims. See *In re Van Geuns*, 988 F.2d 1181, 26 USPQ2d 1057 (Fed. Cir. 1993). Furthermore, Examiner has shown that Leigh's system is a scalable system related to load-balancing. As such, scalable/load-balancing components connected via load-balancing cables are indeed scalability

cables; i.e. scalable/load-balancing components connected via cables are indeed scalability/load-balancing cables. Similarly, load-balancing ports, load-balancing logics of Leigh are indeed *scalability ports* and *scalability logic* respectively. Furthermore, Leigh states that cables could be not Ethernet cables such as USB or 3GIO (paragraph 0058).

As such Examiner respectfully disagrees with Appellant's assertion that *because Leigh does not disclose, anything regarding the cables used to connect Leigh's servers, Leigh cannot disclose scalability cables.*

Appellant argues (page 10, paragraph two) that *Leigh does not disclose scalability ports. Scalability ports, as claimed in the present application, receive scalability cables connecting two systems. As discussed above, Leigh does not disclose scalability cables. Because Leigh does not disclose scalability cables, Leigh cannot possibly disclose scalability ports that receive scalability cables, as claimed here.*

However, Examiner respectfully disagrees with the Appellant's assertion. Examiner has shown above that, Leigh does indeed teach scalability cables. As such ports connected to scalability cables are indeed scalability ports. Specifically, Leigh teaches **(see paragraph 40, the configuration shown is only one possible arrangement of connections among the ports; many other permutations are possible. Not all of the possible connection configurations are optimal for a chosen load balancing method. Leigh in paragraph 40 states, the use of 3-port ILB modules (i.e. A Set Of Scalability Ports) allows the creation (i.e. connoting) of load-balancing clusters (i.e. Scalability Logic) at low cost, since no additional hardware, such**

as a network switch, is needed. This is important especially when there are only a few servers to be load balanced and an ILB module in each host server can provide the load-balanced cluster solution. When there are only a few servers, the installation of the interconnecting cables (i.e. Scalability Cables) among the ILBs is not a significant issue. FIG. 1(d) shows an example of how five ILBs 10, 20, 30, 40, and 50 can be interconnected among themselves and to two external networks via network segments 80 and 90 (i.e. A Set Of Scalability Ports Connected To The Scalability Logic To Receive Scalability Cables Connecting The First System To The Second System). Six inter-ILB connections 83, 84, 85, 86, 87, and 88 are used to interconnect the five 3-port ILBs (i.e. A Set Of Scalability Ports Connected To The Scalability Logic To Receive Scalability Cables Connecting The First System To The Second System). The configuration shown is only one possible arrangement of connections among the ports; many other permutations are possible. Not all of the possible connection configurations are optimal for a chosen load balancing method. Similarly, in the embodiment shown in FIG. 2, twelve 3-port ILBs 202 through 224 are interconnected by sixteen connections and four primary master ILBs 202, 208, 214, and 220 are connected to four external networks via connections 230, 240, 250, and 260 (i.e. A Set Of Scalability Ports Connected To The Scalability Logic To Receive Scalability Cables Connecting The First System To The Second System)) connected to the scalability logic to receive scalability cables (**see paragraph 40 interconnecting cables among the ILBs**) connecting the first system to the second system (**see paragraph 29 and figure 1a-d**).

As such Examiner respectfully disagrees with Appellant's assertion that *because Leigh does not disclose scalability cables or scalability ports, Leigh cannot disclose a set of scalability ports connected to the scalability logic to receive scalability cables connecting the first system to the second system as claimed in the present application.* Specifically, Appellant do not clearly point out the patentable novelty of "Scalability Port" which he or she thinks the claims present in view of the state of the art disclosed by the references cited which teaches "Load-balancing" ports. Examiner has clearly shown above that "load-balancing" is indeed "Scaling". As such, "Scalability Port" is indeed "Load-balancing port". It is for the same reason, Examiner respectfully disagrees with the Appellant's argument (see page 10 last paragraph and 11 first paragraph) that *because Leigh does not disclose scalability cables, scalability ports, or scalability logic. Leigh cannot disclose a set of scalability ports connected to the scalability logic to receive scalability cables connecting tile first system to the second system as claimed in the present application. Because Leigh does not disclose each and every element, and limitation of Appellants' claims, Leigh does not anticipate Appellants' claims, and the rejections tinder 35 U.S.C. § 102 should be withdrawn.*

Appellant argues (page 11, paragraph two) that *Leigh Does Not Disclose System Management To Cause Each Of The System's Scalability Ports to Issue An Identifiable Signal And Further Configured To Detect The Reception Of An Identifiable Signal As Claimed In The Present Application.*

However, Examiner respectfully disagrees with the Appellant's assertion. Leigh does indeed teach the cited limitations. Specifically, Leigh teaches system management

to cause each of the system's scalability ports **(see paragraph 40 the configuration shown is only one possible arrangement of connections among the ports; many other permutations are possible. Not all of the possible connection configurations are optimal for a chosen load balancing method)** to issue an identifiable signal **(see paragraph 30 request)** and further configured to detect the reception of an identifiable signal **(see paragraph 30 ILB 50 acknowledges the load shedding request from ILB 10)**, sent by another system, by any of the scalability ports **(see paragraph 40 the configuration shown is only one possible arrangement of connections among the ports; many other permutations are possible. Not all of the possible connection configurations are optimal for a chosen load balancing method)**. In paragraph 30, Leigh teaches, as an example of zone boundary crossing, in FIG. 1(c), after primary master ILB 10 has distributed the incoming load from network segment 80 to a predefined saturation level on the host servers associated with ILB 10 and ILB 20 in zone-1, ILB 10 (i.e. System Management) may interrogate (i.e. Issue An Identifiable Signal) the other primary master, ILB 50 in zone-2, to possibly accept future incoming loads on network segment 80. If ILB 50 acknowledges (i.e. Reception Of An Identifiable Signal) the load shedding request from ILB 10, then ILB 50 provides ILB 10 with a list of ILBs that can accept the load and with the load-shedding conditions. The load-shedding conditions may include information such as absolute time interval, load-shed check timer value, and number of loads. An example of "absolute time interval" is a wall-clock time interval preset by system administrators based on the known load condition, such as during the first two hours of every weekday, or every business day at

lunch hours. An example of "load-shed check timer value" is a number of hours or minutes, chosen by system administrators, to be set in a register and counted down. This timer value can be set at the time when a load reaches a predefined load index threshold or it can be set periodically. When the timer value has counted down, the corresponding ILB will check the load condition with respect to its resource capacity to determine whether or not it should notify the primary master ILB (if it is not one itself) to redirect the load to another ILB. If ILB 10 sheds its load while complying with the load-shedding conditions, it will stop the load-shedding activity upon a load-shed-abort signal from ILB 50. ILB 50 may issue this load-shed-abort signal to ILB 10 when the servers whose loads are being shed fall below a predefined resource saturation threshold. In regards to *System's Scalability Ports* Leigh in paragraph 0040 teaches the use of 3-port ILB modules (i.e. A Set Of Scalability Ports) allows the creation of load-balancing clusters at low cost, since no additional hardware, such as a network switch, is needed. Leigh teaches in paragraph 0047, at the beginning of the topology discovery algorithm, each ILB will command its individual ports to broadcast a data packet known as a roll-call-1-query packet (i.e. Cause Each Of The System's Scalability Ports to Issue An Identifiable Signal). Each ILB port will respond to the roll-call-1-query packet (i.e. after the Detect of The Reception Of An Identifiable Signal) with a roll-call-1-response packet that includes information such as its port address, the associated ILB ID, and the associated host server ID. Finally, Leigh further teaches in paragraph 0057, when the workload index is consistently low for all the servers, some servers may be selected to be in "sleep" mode to reduce facility power consumption. If the primary master ILB

detects a low workload index for all the servers, it sends a roll-call-sleep command (i.e. Identifiable Signal) packet to the ILBs attached to the servers selected to be in "sleep" mode. After the primary master ILB sends the roll-call-sleep packet to a server, it will not distribute any loads to that server. The primary master ILB adjusts the workload index ratio among the remaining servers accordingly, as described in the adaptive load balancing method. Upon receiving a roll-call-sleep command packet (i.e. Identifiable Signal), a server will respond to the primary master ILB with a roll-call-sleep-wait packet.

As such, Examiner respectfully disagrees with the Appellant's assertion that *Neither Leigh's permutations of arrangements of connections among ports of ILBs nor Leigh's topology detection algorithm that broadcasts a packet and receives responsive packet disclose system management with scalability ports, as claimed in the present application*. Specifically, In response to applicant's argument that the references fail to show certain features of applicant's invention, it is noted that the features upon which applicant relies (i.e., "system management with scalability ports") are not recited in the rejected claim(s). Although the claims are interpreted in light of the specification, limitations from the specification are not read into the claims. See *In re Van Geuns*, 988 F.2d 1181, 26 USPQ2d 1057 (Fed. Cir. 1993).

Since, Leigh indeed teaches Scalability Ports, Examiner respectfully disagrees with the Appellant's assertion that (page 12 last paragraph and 13 first paragraph) *without disclosing a scalability port, Leigh cannot disclose system management using scalability ports as claimed*.

Therefore, claim 1 is deemed not allowable.

35 USC § 102(b) - Claim 8

Appellant argues (page 14 first paragraph) that *Leigh Does Not Disclose Driving An Identifiable Signal On A First Scalability Port Of A First System*.

However, Examiner respectfully disagrees with the Appellant's assertion. Specifically, Leigh does indeed teach the cited limitations. Specifically, Leigh teaches driving an identifiable signal on a first scalability port (**see paragraph 40 the configuration shown is only one possible arrangement of connections among the ports; many other permutations are possible. Not all of the possible connection configurations are optimal for a chosen load balancing method.**) of a first system (**see paragraph 30 request**). Leigh in paragraph 40 states, the use of 3-port ILB modules (*i.e. A Set Of Scalability Ports*) allows the creation (*i.e. connoting*) of load-balancing clusters (*i.e. Scalability Logic*) at low cost, since no additional hardware, such as a network switch, is needed. This is important especially when there are only a few servers to be load balanced and an ILB module in each host server can provide the load-balanced cluster solution. When there are only a few servers, the installation of the interconnecting cables (*i.e. Scalability Cables*) among the ILBs is not a significant issue. FIG. 1(d) shows an example of how five ILBs 10, 20, 30, 40, and 50 can be interconnected among themselves and to two external networks via network segments 80 and 90 (*i.e. A Set Of Scalability Ports Connected To The Scalability Logic To Receive Scalability Cables Connecting The First System To The Second System*). Six

inter-ILB connections 83, 84, 85, 86, 87, and 88 are used to interconnect the five 3-port ILBs (i.e. A Set Of Scalability Ports Connected To The Scalability Logic To Receive Scalability Cables Connecting The First System To The Second System). The configuration shown is only one possible arrangement of connections among the ports; many other permutations are possible. Not all of the possible connection configurations are optimal for a chosen load balancing method. Similarly, in the embodiment shown in FIG. 2, twelve 3-port ILBs 202 through 224 are interconnected by sixteen connections and four primary master ILBs 202, 208, 214, and 220 are connected to four external networks via connections 230, 240, 250, and 260 (i.e. A Set Of Scalability Ports Connected To The Scalability Logic To Receive Scalability Cables Connecting The First System To The Second System)) connected to the scalability logic to receive scalability cables (see paragraph 40 interconnecting cables among the ILBs) connecting the first system to the second system (see paragraph 29 and figure 1a-d). In paragraph 30, Leigh teaches, as an example of zone boundary crossing, in FIG. 1(c), after primary master ILB 10 has distributed the incoming load from network segment 80 to a predefined saturation level on the host servers associated with ILB 10 and ILB 20 in zone-1, ILB 10 (i.e. System Management) may interrogate (i.e. Issue An Identifiable Signal) the other primary master, ILB 50 in zone-2, to possibly accept future incoming loads on network segment 80. If ILB 50 acknowledges (i.e. Reception Of An Identifiable Signal) the load shedding request from ILB 10, then ILB 50 provides ILB 10 with a list of ILBs that can accept the load and with the load-shedding conditions. In regards to *System's Scalability Ports* Leigh in paragraph 0040 teaches the use of 3-port ILB

modules (i.e. A Set Of Scalability Ports) allows the creation of load-balancing clusters at low cost, since no additional hardware, such as a network switch, is needed. Leigh teaches in paragraph 0047, at the beginning of the topology discovery algorithm, each ILB will command its individual ports to broadcast a data packet known as a roll-call-1-query packet (i.e. Cause Each Of The System's Scalability Ports to Issue An Identifiable Signal). Each ILB port will respond to the roll-call-1-query packet (i.e. after the Detect of The Reception Of An Identifiable Signal) with a roll-call-1-response packet that includes information such as its port address, the associated ILB ID, and the associated host server ID. Finally, Leigh further teaches in paragraph 0057, when the workload index is consistently low for all the servers, some servers may be selected to be in "sleep" mode to reduce facility power consumption. If the primary master ILB detects a low workload index for all the servers, it sends a roll-call-sleep command (i.e. Identifiable Signal) packet to the ILBs attached to the servers selected to be in "sleep" mode. After the primary master ILB sends the roll-call-sleep packet to a server, it will not distribute any loads to that server. The primary master ILB adjusts the workload index ratio among the remaining servers accordingly, as described in the adaptive load balancing method. Upon receiving a roll-call-sleep command packet (i.e. Identifiable Signal), a server will respond to the primary master ILB with a roll-call-sleep-wait packet.

As such, Examiner respectfully disagrees with the Appellant's assertion (paragraph one of page 15) that *because Leigh does not disclose scalability cables, Leigh cannot possibly disclose scalability ports that receive scalability cables, as claimed in the present application. Because Leigh does not disclose scalability ports as*

claimed in the present application, Leigh cannot possibly disclose driving an identifiable signal on a first scalability port of a first system as claimed here.

Therefore, claim 8 is deemed not allowable.

It is for the same reasons, contrary to Appellant's assertion (page 15 paragraph 3), Examiner respectfully submits that Leigh does indeed enable each and every element of the claims 1 and 8 of the present Application.

35 USC § 102(b) - Claim 15

Regarding claims 15, Leigh teaches programming (see paragraph 42, **programming**, abstract, A system for balancing network loads among a group of computer servers, paragraph 0003, The present invention generally relates to computer networks and servers and more particularly to a method for balancing network loads among multiple servers, paragraph 0007, A system for balancing external network loads among a group of computer servers includes a load balancing module attached to or integrated within each server), and disclose all the limitations as discussed in the rejection of claims 1 and 8 and are therefore claim 15 is rejected using the same rationales.

As such claim 15 is also deemed not allowable for the same reasons cited above.

35 USC § 102(b) - Claims 2, 9-10 and 16-17

Dependent claims 2, 9-10 and 16-17 are not deemed allowable for the same reasons cited above.

35 USC § 103(a) - Claims 3, 6-7, 11-12, 14, 18-19 and 21

Dependent claims 3, 6-7, 11-12, 14, 18-19 and 21 depend on rejected Independent claims 1, 8 and 15, and are therefore deemed not allowable for the same reasons cited above.

35 USC § 103(a) - Claim 4

Dependent claim 4 depends on rejected Independent claim 1, and is therefore deemed not allowable for the same reasons cited above.

35 USC § 103(a) - Claims 5, 13 and 20

Dependent claims 5, 13 and 20 depend on rejected Independent claims 1, 8 and 15, and are therefore deemed not allowable for the same reasons cited above.

For the above reasons, it is believed that the rejections should be sustained.

(11) Evidence Appendix

No evidence was provided by appellant.

(12) Related Proceeding(s) Appendix

No decision rendered by a court or the Board is identified by the examiner in the Related Appeals and Interferences section of this examiner's answer.

Art Unit: 2419

For the above reasons, it is believed that the rejections should be sustained.

Respectfully submitted,

/Salman Ahmed/

Examiner, Art Unit 2419

/Edan Orgad/

Supervisory Patent Examiner, Art Unit 2419

/Daniel J. Ryman/

Supervisory Patent Examiner, Art Unit 2419